Energy-Efficient Scheduling for Homogeneous Multiprocessor Systems

Mason Thammawichai, *Student Member, IEEE* and Eric C. Kerrigan, *Member, IEEE*

Abstract—We present a number of novel algorithms, based on mathematical optimization formulations, in order to solve a homogeneous multiprocessor scheduling problem, while minimizing the total energy consumption. In particular, for a system with a discrete speed set, we propose solving a tractable linear program. Our formulations are based on a fluid model and a global scheduling scheme, i.e. tasks are allowed to migrate between processors. The new methods are compared with three global energy/feasibility optimal workload allocation formulations. Simulation results illustrate that our methods achieve both feasibility and energy optimality and outperform existing methods for constrained deadline tasksets. Specifically, the results provided by our algorithm can achieve up to an 80% saving compared to an algorithm without a frequency scaling scheme and up to 70% saving compared to a constant frequency scaling scheme for some simulated tasksets. Another benefit is that our algorithms can solve the scheduling problem in one step instead of using a recursive scheme. Moreover, our formulations can solve a more general class of scheduling problems, i.e. any periodic real-time taskset with arbitrary deadline. Lastly, our algorithms can be applied to both online and offline scheduling schemes.

arXiv:1510.05567v2 [cs.OS] 12 Nov 2015

1 Introduction

Due to higher computational power demands in modern computing systems, e.g. sensor networks, satellites, multi-robot systems, as well as personal embedded electronic devices, a system's power consumption is significantly increase. Therefore, an efficient energy management protocol is required in order to balance the power consumption and workload requirement of a system. With this motivation, this paper is aimed at using an optimization-based approach to develop an algorithm for a multiprocessor real-time scheduling problem with the goal of minimizing energy consumption.

The dynamic power consumption of a CMOS processor is often represented as a function of both clock frequency and supply voltage [1]. Thus, a dynamic voltage and frequency scaling (DVFS) scheme is often applied to reduce processor power consumption. DVFS is now commonly implemented in various computing systems, both at hardware and software levels. Examples include Intel's SpeedStep technology, AMD's PowerNow! technology and the Linux Kernel Power Management Scheme.

1.1 Terminologies and Definitions

This section provides basic terminologies and definitions used throughout the paper.

Speed s: The speed s of a processor is defined as the ratio between the operating frequency f and the maximum system frequency f_{max}, i.e. $s := f/f_{max}$. We denote with s_{min} the minimum execution speed of a processor.

Task T_i: A task T_i is defined as a triple $T_i := (c_i, d_i, p_i)$; c_i is the required number of processor cycles, d_i is the task's deadline and p_i is the task's period. If the task's deadline is equal to its period, the task is said to have an 'implicit deadline'. The task is considered to have a 'constrained deadline' if its deadline is not larger than its period, i.e. $d_i \leq p_i$. In the case that the task's deadline can be less than, equal to, or greater than its period, it is said to have an 'arbitrary deadline'.

Job T_{ij}: A job T_{ij} is defined as the j^{th} instance of task T_i, where $j \geq 1$. T_{ij} arrives at time $(j-1)p_i$, has the required execution cycles c_i and a deadline at time $(j-1)p_i + d_i$.

Taskset: The taskset is defined as a set composed of real-time tasks.

Minimum Execution Time $\underline{x}_i$: The minimum execution time $\underline{x}_i$ of a task T_i is the execution time of the task when the task is executed at f_{max}, $\underline{x}_i := c_i/f_{max}$.

Task Density $\delta_i(s_i)$: The task density $\delta_i(s_i)$ of a task T_i executed at a speed s_i is defined as the ratio between the task execution time and the minimum of its deadline and its period, i.e. $\delta_i(s_i) := c_i/(s_i f_{max} \min\{d_i, p_i\})$. When all tasks are assumed to have an implicit deadline, this is often called 'task utilization'.

Taskset Density $D(s_i)$: The taskset density $D(s_i)$ is defined as the summation of all task densities in the taskset, i.e. $D(s_i) := \sum_{i=1}^{n} \delta_i(s_i)$, where n is a number of tasks in the taskset. The minimum taskset density D is given by $D := \sum_{i=1}^{n} \delta_i(1)$.

Feasibility Optimal: Feasibility optimality of a homogeneous multiprocessor system is obtained when the upper bound on the minimum taskset density D, for which the algorithm is able to construct a valid schedule such that no deadlines are missed, equals the total number of processors in the system.

Scheduling Scheme: Multiprocessor scheduling can be classified according to the task migration scheme[1]: a 'global

- *Mason Thammawichai is with the Department of Aeronautics, Imperial College London, London SW7 2AZ, UK*
  E-mail: m.thammawichai12@imperial.ac.uk
- *Eric C. Kerrigan is with the Department of Electrical & Electronic Engineering and the Department of Aeronautics, Imperial College London, London SW7 2AZ, UK.*
  E-mail: e.kerrigan@imperial.ac.uk

1. A task *migration* occurs when a task execution is suspended on one processor and continues on another processor, while *pre-emption* is used when the execution of a task on a processor is suspended in order to start executing another task.

scheduling scheme' allows task migration between processors and a 'partitioned scheduling scheme' does not allow task migration.

1.2 Related Work

There are at least two well-known feasibility optimal homogeneous multiprocessor scheduling algorithms of an implicit deadline taskset that are based on a fluid scheduling model: Proportionate-fair (Pfair) [2] and Largest Local Remaining Execution Time First (LLREF) [3]. Both Pfair and LLREF are global scheduling algorithms. A fluid model, shown in Figure 1, is the ideal schedule path of a task T_i, where the remaining minimum execution time is represented by a straight line and the slope of the line is the task execution speed s_i.

By introducing the notion of *fairness*, Pfair ensures that at any instant no task is one or more quanta (time intervals) away from the task's fluid path. However, the Pfair algorithm suffers from a significant run-time overhead, because tasks are split into several segments, incurring frequent algorithm invocations and task migrations.

To overcome these disadvantages, the LLREF algorithm pre-empts a task at two scheduling events within each time interval [3]. One occurs when the remaining time of an executing task is zero and the other event happens when the difference between a deadline and a remaining execution time of a task is zero.

By incorporating a DVFS scheme with a fluid model, [4] proposed the real-time static voltage and frequency scaling (RT-SVFS) algorithm, which allows the slope of a fluid schedule to vary between 0 and 1. To improve the performance of the open-loop static algorithm of [4], a closed-loop dynamic algorithm is proposed in [5]. Their dynamic algorithm considered uncertainty in the task execution time. By extending [4], an energy-efficient scheduling algorithm for real-time tasks with unknown arrival time, known as a sporadic real-time task, was proposed in [6].

Deadline Partitioning (DP) [7] is the technique that partitions time into intervals bounded by two successive task deadlines, after which each task is allocated the workload and is scheduled at each time interval. A simple optimal scheduling algorithm, called DP-WRAP, was presented in [7]. The DP-WRAP algorithm partitions time according to the DP technique and, at each time interval, the tasks are scheduled using McNaughton's wrap around algorithm [8]. McNaughton's wrap around algorithm aligns all task workloads along a number line, starting at zero, then splits tasks into chunks of length 1 and assigns each chunk to the same processor.

However, the algorithms that are based on the fairness notion [3]–[5], [7], [9], [10] are schedulability optimal, but have hardly been applied in a real system, since they suffer from high scheduling overheads, i.e. task pre-emptions and migrations. Recently, two schedulability optimal algorithms that are not based on the notion of fairness have been proposed. One is the RUN algorithm [11], which uses a dualization technique to reduce the multiprocessor scheduling problem to a series of uniprocessor scheduling problems. The other is U-EDF [12], which generalises the earliest deadline first (EDF) algorithm to multiprocessors by reducing the problem to EDF on a uniprocessor.

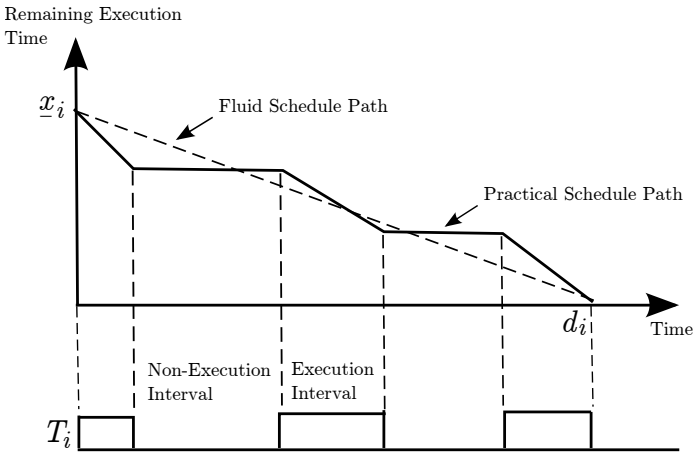

Fig. 1. Fluid Schedule and a Practical Schedule [3]

Alternatively to the above methods, the multiprocessor scheduling problem can also be formulated as an optimization problem. However, since the problem is NP-hard, in general [13], polynomial-time heuristic methods are often used. An example of these approaches can be found in [14], which addresses the energy-aware multiprocessor partitioning problem and [15], which proposes an energy and feasibility optimal global framework for a two-type multicore platform. In general, the tasks are partitioned among the set of processors, followed with computing the running frequency. Among all of the feasibility assignments, an optimal energy consumption assignment is chosen by solving a mathematical optimization problem, where the objective is to minimize energy. The constraints are to ensure that all tasks will meet their deadlines and only one processor is assigned to a task. In partitioned scheduling algorithms, such as [14], once a task is assigned to a specific processor, the multiprocessor scheduling problem is reduced to a set of uniprocessor scheduling problems, which is well studied [16]. However, a partitioned scheduling method cannot provide an optimal schedule.

1.3 Contribution

In this paper, we propose three mathematical optimization problems to solve a periodic hard real-time task scheduling problem on homogeneous multiprocessor systems with DVFS capabilities. First is an MINLP, which adopts the fluid model used in [2]–[4] to represent a scheduling dynamic. The MINLP relies on the optimal control of a suitably-defined dynamic to globally solve for a valid schedule, while solutions are obtained by solving each instance of the time interval obtained using the DP technique [3], [4], [7]. By determining a fraction of task's execution time and operating speed, rather than task assignments, the same scheduling problem can be formulated as an NLP. Lastly, we propose an LP for a system with discrete speed levels. Our work presents homogeneous multiprocessor scheduling algorithms that are both feasibility optimal and energy optimal. Furthermore, our formulations are capable of solving any periodic tasksets, i.e. implict, constrained and arbitary deadlines.

1.4 Outline of Paper

This paper is organized as follows: Section 2 defines our scheduling problem in detail. Three mathematical optimization formulations to solve the same multiprocessor scheduling problem are proposed in Section 3. The simulation setup and results are presented in Section 4. Finally, conclusions and future work are discussed in Section 5.

2 Problem Formulation

2.1 Task and Processor Model

We consider a set of n periodic real-time tasks that are to be partitioned on m identical processors, where each processor's voltage/speed can be adjusted individually. All tasks are assumed to start at the same time. The hyperperiod L is defined as the least common multiple of all task periods. The tasks can be pre-empted at any time, do not share resources and do not have any precedence constraints. It is assumed that $D \leq m$ in order to guarantee the existence of a valid schedule.

Below, we will refer to the sets $I := \{1, \ldots, n\}$, $J_i := \{1, \ldots, L/p_i\}$ and $K := \{1, \ldots, m\}$. The remaining minimum execution time of job T_{ij} at time t will be denoted by $x_{ij}(t)$. Note that $\forall i, \forall j, \forall k, \forall t$ will be used as short-hand for $\forall i \in I, \forall j \in J_i, \forall k \in K, \forall t \in [0, L]$, respectively.

2.2 Energy Consumption Model

For CMOS processors, the total power consumption is often simply expressed as an increasing function of the form

$$P(s) := P_d(s) + P_s = \alpha s^\beta + P_s, \tag{1}$$

where $P_d(s)$ is the dynamic power consumption due to the charging and discharging of CMOS gates, $\alpha > 0$ and $\beta \geq 1$ are hardware-dependent constants and the static power consumption P_s, which is mostly due to leakage current, is assumed to be either constant or zero [1].

The total energy consumption from executing a task can be expressed as a summation of the active and idle energy consumed, i.e. $E = E_{active} + E_{idle}$, where E_{active} is the energy consumed when the processor is busy executing the task and E_{idle} is the energy consumed when the processor is idle. The energy consumed by executing and completing a task at a constant speed s is

$$\begin{aligned} E(s) &= E_{active}(s) + E_{idle}(s) && \text{(2a)} \\ &= \frac{c}{f_{max}} \frac{(P_{active}(s) - P_{idle})}{s} + P_{idle} d && \text{(2b)} \\ &= \frac{x(P_{active}(s) - P_{idle})}{s} + P_{idle} d, && \text{(2c)} \end{aligned}$$

where $P_{active}(s) := P_d(s) + P_s$ is the power consumption in the active interval, P_{idle} is the power consumption during the idle period. $P_{idle} > 0$ and $P_s > 0$ will be assumed to be constants. Note that $P_{active}(s) - P_{idle}$ is strictly greater than zero and that the last term $P_{idle} d$ is not a function of the speed.

2.3 Scheduling as an Optimal Control Problem

The objective is to minimize the total energy consumption of executing a periodic taskset within the hyperperiod L. The scheduling problem can therefore be formulated as the following infinite-dimensional continuous-time optimal control problem:

$$\underset{\substack{x_{ij}(\cdot), a_{ijk}(\cdot), \\ s_{ijk}(\cdot), i \in I, \\ j \in J_i, k \in K}}{\text{minimize}} \sum_{i,j,k} \int_0^L \ell(a_{ijk}(t), s_{ijk}(t)) dt \tag{3a}$$

subject to

$$\begin{aligned} & x_{ij}((j-1)p_i) = \underline{x}_i, && \forall i, j && \text{(3b)} \\ & x_{ij}((j-1)p_i + d_i) = 0, && \forall i, j && \text{(3c)} \\ & \dot{x}_{ij}(t) = -\sum_k s_{ijk}(t) a_{ijk}(t), && \forall i, j, t, \text{ a.e.} && \text{(3d)} \\ & \sum_k a_{ijk}(t) \leq 1, && \forall i, j, t && \text{(3e)} \\ & \sum_{i,j} a_{ijk}(t) \leq 1, && \forall k, t && \text{(3f)} \\ & s_{min} \leq s_{ijk}(t) \leq 1, && \forall i, j, k, t && \text{(3g)} \\ & a_{ijk}(t) \in \{0, 1\}, && \forall i, j, k, t && \text{(3h)} \end{aligned}$$

where

$$\ell(a, s) := a(P_{active}(s) - P_{idle}),$$

$s_{ijk}(t)$ is the execution speed of job T_{ij} at time t and $a_{ijk}(t)$ is used to indicate processor assignment, i.e. $a_{ijk}(t) = 1$ if and only if job T_{ij} is active on processor k at time t.

The initial conditions on the minimum execution time of all jobs are specified in (3b) and job deadline constraints are specified by (3c). The fluid model of the scheduling dynamic is given in (3d), where the state is x and the control inputs are a and s. The constraint that each job is assigned to at most one processor at a time is ensured by (3e) and (3f) enforces that each processor is assigned to at most one job at a time. Upper and lower bounds on the processor speed are given in (3e). The binary nature of job assignment variables are given by (3h).

3 Solving the Scheduling Problem with Finite-dimensional Mathematical Optimization

This section provides details on three mathematical optimization problems to solve the same scheduling problem defined in Section 2.3. The original problem (3) will be discretized by introducing piecewise constant constraints on the control inputs s and a.

Let $\{\tau_0, \tau_1, \ldots, \tau_N\}$, which we will refer to as the major grid, be the set of time instances corresponding to the distinct arrival times and deadlines of all jobs within the hyperperiod L, where $0 = \tau_0 \leq \tau_1 \leq \cdots \leq \tau_N = L$.

3.1 Mixed-integer Nonlinear Program (MINLP-DVFS)

The above scheduling problem, subject to piecewise constant constraints on the control inputs, is most naturally formulated as an MINLP, as defined below.

Though it has been observed by [3] that it is often the case that tasks might have to be split in order to make the entire taskset schedulable, the runtime overheads caused by context switches can jeopardize the performance. Therefore, a variable discretization time step [17] method is applied in a so-called minor grid, so that the solution to our scheduling problem does not depend on the size of the discretization time step. Let $\{\tau_{\mu,0},\ldots,\tau_{\mu,M}\}$ denote the set of time instances on a minor grid within the time interval $[\tau_\mu,\tau_{\mu+1}]$ with $\tau_\mu=\tau_{\mu,0}\leq\cdots\leq\tau_{\mu,M}=\tau_{\mu+1}$, so that $\{\tau_{\mu,1},\ldots,\tau_{\mu,M-1}\}$ is to be determined for all μ from solving an appropriately-defined optimization problem.

Let $\forall\mu$ and $\forall\nu$ be short-hand for $\forall\mu\in U:=\{0,\ldots,N-1\}$ and $\forall\nu\in V:=\{0,\ldots,M-1\}$. Define the notation $[\mu,\nu]:=(\tau_{\mu,\nu}),\forall\mu,\nu$ and the discretized state and input sequences as

$$\begin{aligned}
x_{ij}[\mu,\nu]&:=x_{ij}(\tau_{\mu,\nu}), && \forall i,j,\mu,\nu\\
s_{ijk}[\mu,\nu]&:=s_{ijk}(\tau_{\mu,\nu}), && \forall i,j,k,\mu,\nu\\
a_{ijk}[\mu,\nu]&:=a_{ijk}(\tau_{\mu,\nu}), && \forall i,j,k,\mu,\nu
\end{aligned}$$

If $s_{ijk}(\cdot)$ and $a_{ijk}(\cdot)$ are constant in-between time instances on the minor grid, i.e.

$$s_{ijk}(t)=s_{ijk}[\mu,\nu],\ \forall t\in[\tau_{\mu,\nu},\tau_{\mu,\nu+1}),\mu,\nu \tag{5a}$$
$$a_{ijk}(t)=a_{ijk}[\mu,\nu],\ \forall t\in[\tau_{\mu,\nu},\tau_{\mu,\nu+1}),\mu,\nu \tag{5b}$$

then it is easy to show that the solution of the scheduling dynamic (3d) is given by

$$\begin{aligned}
x_{ij}[\mu,\nu+1]&=x_{ij}[\mu,\nu]\\
&\quad-h[\mu,\nu]\sum_k s_{ijk}[\mu,\nu]a_{ijk}[\mu,\nu],\quad \forall i,j,\mu,\nu
\end{aligned}\tag{6a}$$

where

$$h[\mu,\nu]:=\tau_{\mu,\nu+1}-\tau_{\mu,\nu},\forall\mu,\nu. \tag{6b}$$

Let Λ denote the set of all jobs within hyperperiod L, i.e. $\Lambda:=\{T_{ij}|i\in I,j\in J_i\}$. Define a function $\Phi_A:\Lambda\to U\times V$ by $\Phi_A(T_{ij}):=(\mu,\nu)$ such that $\tau_{\mu,\nu}=(j-1)p_i,\ \forall T_{ij}\in\Lambda$ and a function $\Phi_D:\Lambda\to U\times V$ by $\Phi_D(T_{ij}):=(\mu,\nu)$ such that $\tau_{\mu,\nu}=(j-1)p_i+d_i,\ \forall T_{ij}\in\Lambda$.

The original problem (3) subject to the piecewise constant constraints on the inputs (5) is therefore equivalent to the following finite-dimensional MINLP:

$$\underset{\substack{x_{ij}[\cdot],s_{ijk}[\cdot],\\ a_{ijk}[\cdot],h[\cdot],\\ i\in I,j\in J_i,k\in K}}{\text{minimize}}\ \sum_{\mu,\nu,i,j,k} h[\mu,\nu]\ell(a_{ijk}[\mu,\nu],s_{ijk}[\mu,\nu]) \tag{6c}$$

subject to (6a) and

$$x_{ij}[\Phi_A(T_{ij})]=\underline{x}_i \qquad \forall i,j \tag{6d}$$
$$x_{ij}[\Phi_D(T_{ij})]=0, \qquad \forall i,j \tag{6e}$$
$$\sum_k a_{ijk}[\mu,\nu]\leq 1, \qquad \forall i,j,\mu,\nu \tag{6f}$$
$$\sum_{i,j} a_{ijk}[\mu,\nu]\leq 1, \qquad \forall k,\mu,\nu \tag{6g}$$
$$s_{min}\leq s_{ijk}[\mu,\nu]\leq 1, \qquad \forall i,j,k,\mu,\nu \tag{6h}$$
$$a_{ijk}[\mu,\nu]\in\{0,1\}, \qquad \forall i,j,k,\mu,\nu \tag{6i}$$
$$0\leq h[\mu,\nu], \qquad \forall\mu,\nu \tag{6j}$$
$$\sum_\nu h[\mu,\nu]=\tau_{\mu+1}-\tau_\mu, \qquad \forall\mu \tag{6k}$$

where (6j)–(6k) enforce upper and lower bounds on the discretization time steps.

Theorem 1. If a solution exists to (3b)–(3h) and $M\geq\max\{n-1,m-1\}$, then a solution exists to problem (6).

Proof: Follows from the fact that if a solution exists to (3b)–(3h), then the DP-WRAP scheduling algorithm can find a valid schedule with at most $n-1$ context switches and $m-1$ migrations per slice [7, Thm 6]. □

3.2 Continuous Nonlinear Program (NLP-DVFS)

This section proposes an NLP formulation without integer variables to solve the problem in Section 2.3. The idea is to relax the binary constraints in (6i) so that the value of a can be interpreted as the fraction of a time interval during which the job is executed on a processor.

Moreover, the number of variables can also be reduced compared to problem (6), since the processor assignment information does not help in finding the task execution order, hence k is dropped from the subscripts in the notation. That is, partitioning time using only the major grid (i.e. $M=1$) is enough to find a feasible schedule if a solution exists to the original problem (3). Since we only need a major grid, we define the notation $[\mu]:=(\tau_\mu)$.

Consider now the following finite-dimensional NLP:

$$\underset{\substack{x_{ij}[\cdot],s_{ij}[\cdot],\\ a_{ij}[\cdot],i\in I,j\in J_i}}{\text{minimize}}\ \sum_{\mu,i,j}(\tau_{\mu+1}-\tau_\mu)\ell(a_{ij}[\mu],s_{ij}[\mu]) \tag{7a}$$

subject to

$$x_{ij}[\Phi_A(T_{ij})]=\underline{x}_i, \qquad \forall i,j \tag{7b}$$
$$x_{ij}[\Phi_D(T_{ij})]=0, \qquad \forall i,j \tag{7c}$$
$$\begin{aligned}
x_{ij}[\mu+1]&=x_{ij}[\mu]\\
&\quad-(\tau_{\mu+1}-\tau_\mu)s_{ij}[\mu]a_{ij}[\mu], \qquad \forall i,j,\mu
\end{aligned}\tag{7d}$$
$$\sum_{i,j} a_{ij}[\mu]\leq m, \qquad \forall\mu \tag{7e}$$
$$s_{min}\leq s_{ij}[\mu]\leq 1, \qquad \forall i,j,\mu \tag{7f}$$
$$0\leq a_{ij}[\mu]\leq 1, \qquad \forall i,j,\mu \tag{7g}$$

where $a_{ij}[\mu]\in\mathbb{R}$ is the fraction of the time interval $[\tau_\mu,\tau_{\mu+1})$ for which job T_{ij} is executing on a processor at speed $s_{ij}[\mu]$ and (7e) specifies that the total workload in time interval $[\tau_\mu,\tau_{\mu+1})$ should be less than or equal to the system capacity.

Theorem 2. If a solution exists to (3b)–(3h), then a solution to problem (7) exists. Furthermore, at least one valid schedule satisfying (3b)–(3h) can be constructed from a solution to problem (7) with the same energy consumption.

Proof: A similar argument as Theorem 1 can be used to prove the existence of a solution. The simplest valid schedule can be constructed using McNaughton's wrap around algorithm [8] for each time interval $[\tau_\mu,\tau_{\mu+1})$. During each interval $[\tau_\mu,\tau_{\mu+1})$, the taskset is schedulable since (i) the total density of the taskset does not exceed the total number of processors, which is guaranteed by contraint (7e), (ii) no task workload is greater than its deadline, which is ensured by constraint (7g) and (iii) task migration is allowed, which is our assumption. The optimal energy

consumption relies on the fact that the total energy consumption does not depend on the order of which tasks are scheduled, but rather depends on the taskset's density, i.e. the objective value stays the same regardless of the number of discretization steps M on the minor grid. □

3.3 Linear Program (LP-DVFS)

Suppose now that the set of speed levels is finite, as is the case with a real processor. We denote with s_q the speed of the processor at level $q \in Q := \{1, \ldots, l\}$, where l is the total number of speed levels.

Consider now the following finite-dimensional LP:

$$\underset{\substack{x_{ij}[\cdot], a_{ij}^q[\cdot], \\ i\in I, j\in J_i, q\in Q}}{\text{minimize}} \quad \sum_{\mu,i,j,q} (\tau_{\mu+1} - \tau_\mu)\ell(a_{ij}^q[\mu], s_q) \tag{8a}$$

subject to

$$x_{ij}[\Phi_A(T_{ij})] = \underline{x}_i \qquad \forall i,j \tag{8b}$$

$$x_{ij}[\Phi_D(T_{ij})] = 0, \qquad \forall i,j \tag{8c}$$

$$x_{ij}[\mu+1] = x_{ij}[\mu] - (\tau_{\mu+1} - \tau_\mu)\sum_q s_q a_{ij}^q[\mu], \qquad \forall i,j,\mu \tag{8d}$$

$$\sum_q a_{ij}^q[\mu] \le 1, \qquad \forall i,j,\mu \tag{8e}$$

$$\sum_{i,j,q} a_{ij}^q[\mu] \le m, \qquad \forall \mu \tag{8f}$$

$$0 \le a_{ij}^q[\mu] \le 1, \qquad \forall i,j,q,\mu \tag{8g}$$

where $a_{ij}^q[\mu] \in \mathbb{R}$ is the fraction of the time interval $[\tau_\mu, \tau_{\mu+1})$ for which job T_{ij} is executing on a processor at speed s_q.

Constraint (8e) assures that a task will not run on more than one processor at a time. Constraint (8f) guarantees that a processor's workload will not exceed its capacity at a time. Lastly, constraint (8g) provides upper and lower bounds on a fraction of a job's execution time at a specific speed level variable.

Note that, given a solution to the optimization problem (8), one could also employ McNaughton's wrap around algorithm [8] to construct a valid schedule of processor assignments with the same energy consumption. In other words, Theorem 2 can be applied here as well to prove an existence of a valid schedule given a solution to problem (8).

4 Simulation Results

4.1 System, Processor and Task Models

The performance of solving the above optimization problems is evaluated on two models of practical commercial processors, namely an XScale and PowerPC 405LP. The power consumption details of the two commercial processors, which have also been used in [18]–[20], are given in Table 1. The active power consumption models of the XScale and PowerPC 405LP shown in Table 2 were obtained by a polynomial curve fitting to the generic form (1) (details are given in the appendix B). The plots of the actual data versus the fitted models are shown in Fig. 2.

4.2 Comparison between Algorithms

For a system with a continuous speed range, four algorithms were compared: (i) MINLP-DVFS, (ii) NLP-DVFS, (iii) GP-SVFS, which represents a global energy/feasibility-optimal workload partitioning with constant frequency scaling scheme and (iv) GP-NoDVFS, which is a global workload allocation without frequency scaling scheme. For a system with discrete speed levels, three algorithms are compared: (i) LP-DVFS, (ii) GP-NoDVFS and (iii) GP-SDiscrete, which represents global energy/feasibility-optimal workload allocations with constant discrete frequency scaling. Note that the formulations of GP-SVFS and GP-SDiscrete are similar to [14], [15]. Specifically, the GP-SVFS is based on a constant frequency scaling with global scheduling scheme, while [14] is a partitioning-based formulation and [15] is a generalized formulation for a two-type heterogeneous multi-core system. GP-SDiscrete is an extension to systems with discrete speed levels. Details on GP-SVFS, GP-NoDVFS and GP-SDiscrete are given below.

GP-SVFS/GP-NoDVFS: Solve

$$\underset{s_k, k\in K}{\text{minimize}} \sum_{i,k} L\ell(u_{ik}(s_k), s_k) \tag{9a}$$

$$\text{subject to} \sum_k u_{ik}(s_k) \le 1, \quad \forall i \tag{9b}$$

$$\sum_i u_{ik}(s_k) \le 1, \qquad \forall k \tag{9c}$$

$$s_{min} \le s_k \le 1, \qquad \forall k\ (GP-SVFS) \tag{9d}$$

$$s_k = 1, \qquad \forall k\ (GP-NoDVFS) \tag{9e}$$

where $u_{ik}(s_k)$ is the task density on processor k, i.e. $u_{ik}(s_k) := \underline{x}_i/(s_k f_{max} d_i)$ and s_k is the static execution speed of processor k. A task will not be executed on more than one processor at the same time due to (9b) and the assigned workload will not exceed processor capacity due to (9c). The difference between GP-SVFS and GP-NoDVFS lies in restriction on operating speed to be either a continous variable (9d) or fixed at the maximum speed (9e).

GP-SDiscrete: Determine a fraction of the workload y_{ikq} of a task at a specific speed level and a processor speed level selection z_{kq} such that:

$$\underset{\substack{y_{ikq}, z_{kq}, i\in I, \\ k\in K, q\in Q}}{\text{minimize}} \sum_{i,k,q} L\ell(u_{ikq}(y_{ikq})z_{kq}, s_q) \tag{10a}$$

$$\text{subject to} \sum_{k,q} y_{ikq} z_{kq} = 1, \qquad \forall i \tag{10b}$$

$$\sum_q z_{kq} = 1, \qquad \forall k \tag{10c}$$

$$\sum_{k,q} u_{ikq}(y_{ikq})z_{kq} \le 1, \qquad \forall i \tag{10d}$$

$$\sum_{i,q} u_{ikq}(y_{ikq})z_{kq} \le 1, \qquad \forall k \tag{10e}$$

$$0 \le y_{ikq} \le 1, \qquad \forall i,k,q \tag{10f}$$

$$z_{kq} \in \{0,1\}, \qquad \forall k,q \tag{10g}$$

where y_{ikq} represents a fraction of the workload of a task at a specific speed level and z_{kq} is a speed level selection variable for processor k, i.e. $z_{kq} = 1$ if a speed level q is selected and $z_{kq} = 0$ otherwise. The constraints in (10b)

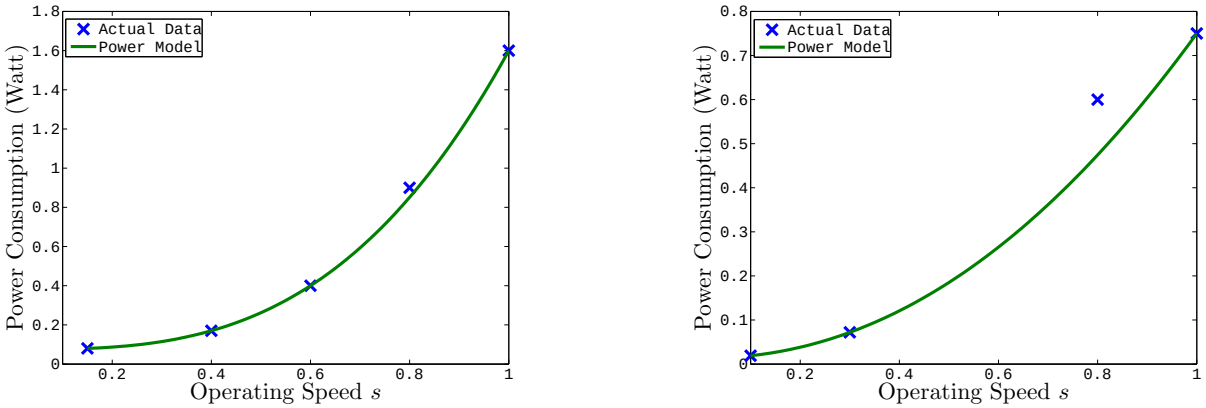

(a) XScale (b) PowerPC 405LP

Fig. 2. Actual Data versus Fitted Model

TABLE 1
Commercial Processor Details for Simulation

Processor type	XScale [21]					PowerPC 405LP [22]			
Frequency (MHz)	150	400	600	800	1000	33	100	266	333
Speed	0.15	0.4	0.6	0.8	1.0	0.1	0.3	0.8	1.0
Voltage (V)	0.75	1.0	1.3	1.6	1.8	1.0	1.0	1.8	1.9
Active Power (mW)	80	170	400	900	1600	19	72	600	750
Idle Power (mW)	40 [18]					12			

TABLE 2
Processor Power Consumption Fitted Model

Processor type	Active Power Model	MAPE[2]
XScale	$P_{active}(s) := 1524.92s^{3.0269} + 75.1092$	1.1236
PowerPC 405LP	$P_{active}(s) := 736.87s^{2.0990} + 13.1333$	5.2323

TABLE 3
Simulation Tasksets

D	T_1	T_2	T_3	T_4
0.4	(0.75,5,10)	(0.75,5,10)	(0.5,10,10)	(0.5,10,10)
0.6	(1,5,10)	(1,5,10)	(1,10,10)	(1,10,10)
0.8	(1.5,5,10)	(1.5,5,10)	(1,10,10)	(1,10,10)
1.0	(2,5,10)	(2,5,10)	(1,10,10)	(1,10,10)
1.2	(2.5,5,10)	(2.5,5,10)	(1,10,10)	(1,10,10)
1.4	(3,5,10)	(3,5,10)	(1,10,10)	(1,10,10)
1.6	(3,5,10)	(3,5,10)	(2,10,10)	(2,10,10)
1.8	(3.5,5,10)	(3.5,5,10)	(2,10,10)	(2,10,10)
2.0	(4,5,10)	(4,5,10)	(2,10,10)	(2,10,10)

Note: The first parameter of a task is $\underline{x}_i$; c_i can be obtained by multiplying $\underline{x}_i$ by f_{max}.

guarantee that the total workload of a task is equal to 1. The constraints in (10c) ensure that only one speed level is selected, (10d) assures that a task will be executed on only one processor at a time, (10f) ensures that a processor workload capacity is not violated and (10g) emphasises that the speed level selection variable is binary. Note that GP-SVFS and GP-NoDVFS are NLPs and GP-SDiscrete is an MINLP.

2. See Appendix A.

4.3 Simulation Setup and Results

For simplicity, we consider the case where four independent periodic real-time tasks with constrained deadlines need to be scheduled onto two homogeneous processors. The total energy consumption of each taskset in Table 3 were evaluated. For the MINLP-DVFS, NLP-DVFS and LP-DVFS implementations, the major grid discretization step $N := 2$ because all tasksets for the simulation only have two distinct deadlines. For the MINLP-DVFS, we chose the minor grid discretization step $M := 4$, since there are at most four jobs in each major grid, which implies that there will be at most four preemptions within each major grid. All of our formulations were modelled using ZIMPL [23] and solved by SCIP [24].

Simulation results are shown in Figures 3 and 4. The minimum taskset density D is represented on the horizontal axis. The vertical axis represents the total energy consumption normalised by the energy used by GP-NoDVFS, i.e. less than 1 means that an algorithm does better than GP-NoDVFS. It can be seen from the plots that for a time-varying workload such as a constrained deadline taskset, our results from solving MINLP-DVFS, NLP-DVFS and LP-DVFS are energy optimal, while GP-SVFS, GP-Discrete and GP-NoDVFS are not. This is because our formulations are incorporated with time, which provides benefits on solving a scheduling problem with both time-varying workload

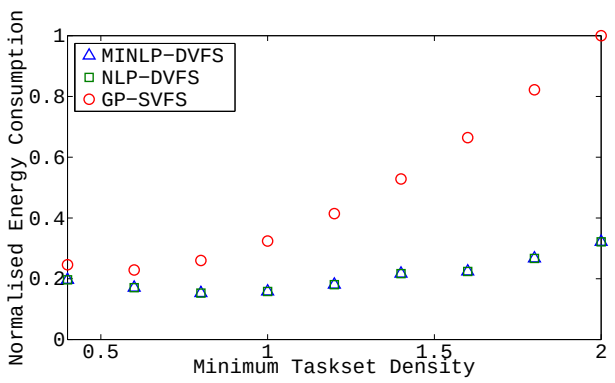

(a) Xscale

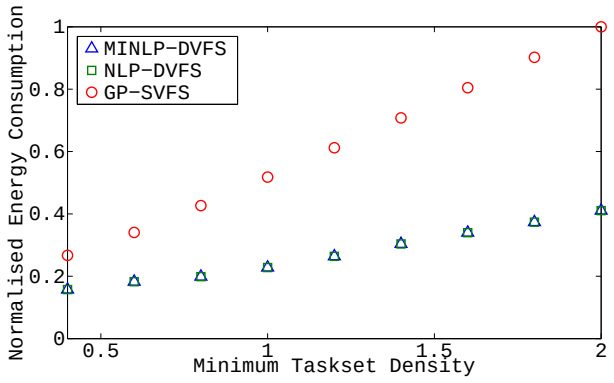

(b) PowerPC 405LP

Fig. 3. Simulation results for scheduling real-time tasks with constrained deadlines on two homogeneous processors with continuous speed range

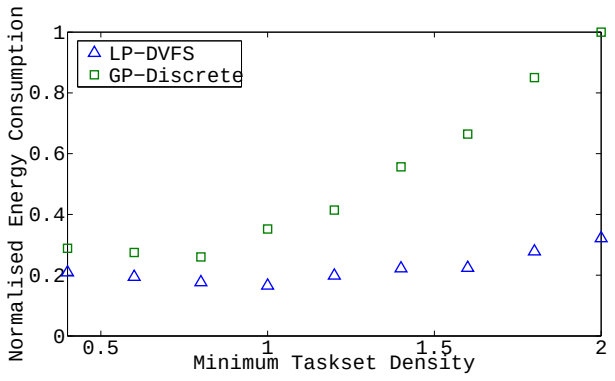

(a) Xscale

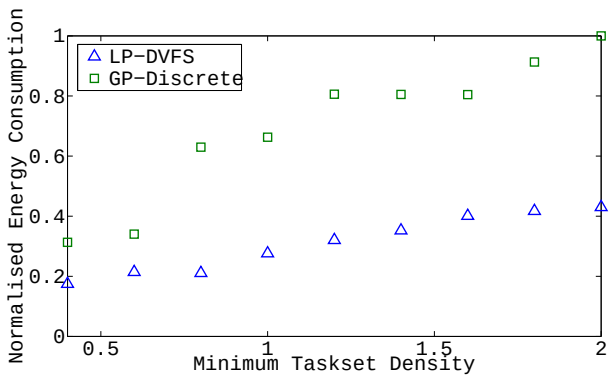

(b) PowerPC 405LP

Fig. 4. Simulation results for scheduling real-time tasks with constrained deadlines on two homogeneous processors with discrete speed levels

(constrained deadline taskset) and constant workload (implicit deadline taskset). In general, compared to the constant speed profile, the time-varying speed profile saving increases as the minimum taskset density increases. It can also be noticed that the percentage saving of the Xscale is nonlinear compared to that of the PowerPC's. This is because the power consumption model of the Xscale is a cubic, while the PowerPC's power consumption model is a quadratic. However, it has to be mentioned that the energy saving percentage varies with the taskset, which implies that the number shown on the plots can be varied, but the significant outcomes stay the same. Lastly, the computation times to solve NLP-DVFS, LP-DVFS, GP-SVFS and GP-SDiscrete are extremely fast, i.e. milliseconds using a general-purpose desktop PC with off-the-shelf optimization solvers, while the time to solve MINLP-DVFS can be up to an hour in some cases.

5 Conclusions

Three mathematical optimization problems were proposed to solve a homogeneous multiprocessor scheduling problem with a periodic real-time taskset. Though our MINLP and NLP formulations are both energy optimal and feasibility optimal, the computation time is high compared with heuristic algorithms. However, our LP formulation is computationally tractable and suitable for a system with discrete speed level set. Moreover, we have shown via simulations that our formulations are able to solve a more general class of scheduling problem than existing work in the literature due to incorporating a scheduling dynamic model in the formulations as well as allowing for a time-varying executing speed profile. The simulation results illustrate the possibility that the time-varying speed profile can save energy up to 70% compared to the constant speed profile. Possible future work could also include developing numerically efficient methods to solve the various mathematical optimization problems defined in this paper. One could also extend the ideas presented here to solve a dynamic scheduling problem with uncertainty in task execution time and to include slack reclamation for further energy reduction.

References

[1] J. M. Rabaey, A. P. Chandrakasan, and B. Nikolic, *Digital integrated circuits : a design perspective*, 2nd ed., ser. Prentice Hall electronics and VLSI series. Pearson Education, Jan. 2003.

[2] S. K. Baruah, N. K. Cohen, C. G. Plaxton, and D. A. Varvel, "Proportionate progress: A notion of fairness in resource allocation," in *Proceedings of the Twenty-fifth Annual ACM Symposium on Theory of Computing*, ser. STOC '93. New York, NY, USA: ACM, 1993, pp. 345–354. [Online]. Available: http://doi.acm.org/10.1145/167088.167194

[3] H. Cho, B. Ravindran, and E. Jensen, "An optimal real-time scheduling algorithm for multiprocessors," in *Real-Time Systems Symposium, 2006. RTSS '06. 27th IEEE International*, Dec 2006, pp. 101–110.

[4] K. Funaoka, S. Kato, and N. Yamasaki, "Energy-efficient optimal real-time scheduling on multiprocessors," in *Object Oriented Real-Time Distributed Computing (ISORC), 2008 11th IEEE International Symposium on*, May 2008, pp. 23–30.

[5] K. Funaoka, A. Takeda, S. Kato, and N. Yamasaki, "Dynamic voltage and frequency scaling for optimal real-time scheduling on multiprocessors," in *Industrial Embedded Systems, 2008. SIES 2008. International Symposium on*, June 2008, pp. 27–33.

[6] D.-S. Zhang, F.-Y. Chen, H.-H. Li, S.-Y. Jin, and D.-K. Guo, "An energy-efficient scheduling algorithm for sporadic real-time tasks in multiprocessor systems," in *High Performance Computing and Communications (HPCC), 2011 IEEE 13th International Conference on*, Sept 2011, pp. 187–194.

[7] G. Levin, S. Funk, C. Sadowski, I. Pye, and S. Brandt, "Dp-fair: A simple model for understanding optimal multiprocessor scheduling," in *Real-Time Systems (ECRTS), 2010 22nd Euromicro Conference on*, July 2010, pp. 3–13.

[8] R. McNaughton, "Scheduling with deadlines and loss function," *Machine Science*, vol. 6(1), pp. 1–12, October 1959.

[9] S. Funk, V. Berten, C. Ho, and J. Goossens, "A global optimal scheduling algorithm for multiprocessor low-power platforms," in *Proceedings of the 20th International Conference on Real-Time and Network Systems*, ser. RTNS '12. New York, NY, USA: ACM, 2012, pp. 71–80. [Online]. Available: http://doi.acm.org/10.1145/2392987.2392996

[10] F. Wu, S. Jin, and Y. Wang, "A simple model for the energy-efficient optimal real-time multiprocessor scheduling," in *Computer Science and Automation Engineering (CSAE), 2012 IEEE International Conference on*, vol. 3, May 2012, pp. 18–21.

[11] P. Regnier, G. Lima, E. Massa, G. Levin, and S. Brandt, "Run: Optimal multiprocessor real-time scheduling via reduction to uniprocessor," in *Real-Time Systems Symposium (RTSS), 2011 IEEE 32nd*, Nov 2011, pp. 104–115.

[12] G. Nelissen, V. Berten, V. Nelis, J. Goossens, and D. Milojevic, "U-edf: An unfair but optimal multiprocessor scheduling algorithm for sporadic tasks," in *Real-Time Systems (ECRTS), 2012 24th Euromicro Conference on*, July 2012, pp. 13–23.

[13] E. Lawler, "Recent results in the theory of machine scheduling," in *Mathematical Programming The State of the Art*, A. Bachem, B. Korte, and M. Grtschel, Eds. Springer Berlin Heidelberg, 1983, pp. 202–234. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-68874-4_9

[14] H. Aydin and Q. Yang, "Energy-aware partitioning for multiprocessor real-time systems," in *Parallel and Distributed Processing Symposium, 2003. Proceedings. International*, April 2003, pp. 9 pp.–.

[15] H. S. Chwa, J. Seo, H. Yoo, J. Lee, and I. Shin, "Energy and feasibility optimal global scheduling framework on big.little platforms," Department of Computer Science, KAIST and Department of Computer Science and Engineering, Sungkyunkwan University, Republic of Korea, Tech. Rep., 2014. [Online]. Available: https://cs.kaist.ac.kr/upload_files/report/1407392146.pdf

[16] J.-J. Chen and C.-F. Kuo, "Energy-efficient scheduling for real-time systems on dynamic voltage scaling (DVS) platforms," in *Embedded and Real-Time Computing Systems and Applications, 2007. RTCSA 2007. 13th IEEE International Conference on*, Aug 2007, pp. 28–38.

[17] M. Gerdts, "A variable time transformation method for mixed-integer optimal control problems," *Optimal Control Applications and Methods*, vol. 27, no. 3, pp. 169–182, 2006. [Online]. Available: http://dx.doi.org/10.1002/oca.778

[18] R. Xu, C. Xi, R. Melhem, and D. Moss, "Practical pace for embedded systems," in *Proceedings of the 4th ACM International Conference on Embedded Software*, ser. EMSOFT '04. New York, NY, USA: ACM, 2004, pp. 54–63. [Online]. Available: http://doi.acm.org/10.1145/1017753.1017767

[19] G. Zeng, T. Yokoyama, H. Tomiyama, and H. Takada, "Practical energy-aware scheduling for real-time multiprocessor systems," in *Embedded and Real-Time Computing Systems and Applications, 2009. RTCSA '09. 15th IEEE International Conference on*, Aug 2009, pp. 383–392.

[20] H.-C. Wang, I. Woungang, C.-W. Yao, A. Anpalagan, and M. S. Obaidat, "Energy-efficient tasks scheduling algorithm for real-time multiprocessor embedded systems," *J. Supercomput.*, vol. 62, no. 2, pp. 967–988, Nov. 2012. [Online]. Available: http://dx.doi.org/10.1007/s11227-012-0771-0

[21] I. X. M. Benchmarks, 2005, http://web.archive.org/web/20050326232506/developer.intel.com/design/intelxscale/benchmarks.htm.

[22] C. Rusu, R. Xu, R. Melhem, and D. Mosse, "Energy-efficient policies for request-driven soft real-time systems," in *Real-Time Systems, 2004. ECRTS 2004. Proceedings. 16th Euromicro Conference on*, June 2004, pp. 175–183.

[23] T. Koch, "Rapid mathematical prototyping," Ph.D. dissertation, Technische Universität Berlin, 2004.

[24] T. Achterberg, "SCIP: Solving constraint integer programs," *Mathematical Programming Computation*, vol. 1, no. 1, pp. 1–41, July 2009, http://mpc.zib.de/index.php/MPC/article/view/4.

APPENDIX A
MEAN ABSOLUTE PERCENTAGE ERROR (MAPE)

$$\text{MAPE} := \frac{1}{k}\sum_{i=1}^{k}\frac{|F(z(i)) - y(i)|}{|y(i)|} \times 100, \tag{11}$$

where $\frac{|F(z(i))-y(i)|}{|y(i)|}$ is the magnitude of the relative error in the i^{th} measurement, $z \mapsto F(z)$ is the estimated function, z is the input data, y is the actual data and k is the total number of fitted points.

APPENDIX B
CURVE FITTING

The following shows the curve fitting formulation to obtain the active power consumption function of the XScale and PowerPC 405LP with the objective of minimizing the MAPE:

$$\underset{\alpha,\beta,P_s}{\text{minimize}}\ 100 \times \frac{1}{\gamma}\sum_{i=1}^{\gamma}|\alpha s_i^{\beta} + P_s - P_{active}(s_i)|/P_{active}(s_i)$$

subject to

$$P_s \geq 0 \tag{12a}$$

$$\alpha \geq 0 \tag{12b}$$

$$\beta \geq 1 \tag{12c}$$

where α, β and P_s are parameters to be fitted, γ is the total number of speed levels of a processor. The speed level of a processor s_i and the active power consumption of a processor $P_{active}(s_i)$ are measurement data shown in Table 1.